

Dynamics of episodic transient correlations in currency exchange rate returns and their predictability

Milan Žukovič^{1,2}

¹ *SORS Research a.s., Moyzesova 38, 040 01 Košice, Slovakia*

² *Department of Theoretical Physics and Astrophysics, Faculty of Science, P. J. Šafárik University, Park Angelinum 9, 041 54 Košice, Slovakia*

Abstract

We study the dynamics of the linear and non-linear serial dependencies in financial time series in a rolling window framework. In particular, we focus on the detection of episodes of statistically significant two- and three-point correlations in the returns of several leading currency exchange rates that could offer some potential for their predictability. We employ a rolling window approach in order to capture the correlation dynamics for different window lengths and analyze the distributions of periods with statistically significant correlations. We find that for sufficiently large window lengths these distributions fit well to power-law behavior. We also measure the predictability itself by a hit rate, i.e. the rate of consistency between the signs of the actual returns and their predictions, obtained from a simple correlation-based predictor. It is found that during these relatively brief periods the returns are predictable to a certain degree and the predictability depends on the selection of the window length.

Keywords: Financial time series, episodic non-linearity, bicornelation, time series prediction

1. Introduction

Recent empirical studies suggested that returns of financial time series depart from the random walk hypothesis by showing a certain degree of

Email address: milan.zukovic@upjs.sk (Milan Žukovič^{1,2})

long-term or short-term dependent relationships, thus violating the weak-form efficient market hypothesis [1, 2, 3]. The efficiency/non-efficiency can, for example, be assessed by evaluation of the Hurst exponent [4] or the approximate entropy [5], as measures of long-term memory and randomness in time series. Since the efficient market hypothesis in its weakest form implies that the returns should be serially uncorrelated, its validity can also be assessed by looking for evidence of significant linear [6, 7] and non-linear [8] serial autocorrelations (henceforth, correlations) in the returns. A systematic review of literature on the weak-form market efficiency of stock markets using a wide array of statistical tests, such as the linear serial correlations, unit root, low-dimensional chaos, nonlinear serial dependence and long memory was provided by Lim and Brooks [9]. However, the market efficiency vary with time and the character of possible serial dependencies is not known. Moreover, they are of transient nature, showing up in random intervals just for a short time, which makes it difficult to exploit them for prediction purposes [10].

In several previous studies [10, 11, 12, 13] the non-linearities were investigated by calculation of the bicorrelation test statistic, due to Hinich [14], and performing the windowed-test procedure [15, 16]. In order to capture the time variation, the data were split into a set of non-overlapping windows, the length of which was set to some ad-hoc fixed value. However, considering the episodic and transient nature of the correlations, these may or may not be detected, depending on the window length used. Furthermore, the window length also influences the onset and offset of the significant correlations. Shorter lengths facilitate quicker response to changes in the correlation strength and can help pinpoint the arrival and disappearance of the transient dependences, but on the other hand, they may lack adequate statistical power.

In order to track the evolution of market efficiency over time, a rolling-window approach has been adopted to calculate a time-varying Hurst exponent [17, 18, 19, 20]. Lim [21] has proposed a similar rolling window approach for computing the bicorrelation test for various stock markets [22, 23]. It was shown that the market efficiency is not a static property that remains unchanged throughout the entire estimation period. Thus, examination of the presence of significant correlations in the rolling sample framework can provide a useful tool for tracking the changing degree of weak-form market efficiency over time as well as for its ranking among different markets. In the currency exchange market, Brooks and Hinich [24] found that Sterling

exchange rates are characterized by transient epochs of dependencies surrounded by long periods of white noise.

The objective of this paper is to study the dynamics of the linear and non-linear serial dependencies in financial time series, using a portmanteau test procedure in a rolling window framework. In particular, we focus on detection of episodes of statistically significant two- and three-point correlations in the returns of several leading currency exchange rates that could offer some potential for their predictability. We employ a rolling window approach in order to capture the correlation dynamics for different window lengths. In comparison with similar previous studies, our new contribution is using tools from statistical physics in order to analyze distributions of the periods with statistically significant correlations and relate them to the predictability obtained from a simple correlation-based predictor.

2. Data

We use hourly average exchange rates per Euro (EUR) on the US dollar (USD), Canadian dollar (CAD), Swiss franc (CHF), British pound (GBP) and Japanese yen (JPY). The sample period is from March 1, 2004 to December 17, 2009 (a total of 37,620 observations). In Fig. 1 we plot the nominal exchange rates of the respective currencies vs the Euro. In Figs. 2 and 3 we show first logarithmic differences of the nominal exchange rates and their frequency distributions, respectively. The histograms and the summary statistics, presented in Table 1, indicate that the data deviate from the normal distribution. This observation is also supported by the Jarque-Bera [25] normality tests, which reject the null hypothesis of normality with very low p -values.

3. Methods

3.1. Detection of significant correlations and bicorrelations in a rolling window

Let $R(t)$ represent the logarithmic returns of the length N , defined by $R(t) = \ln(P(t+1)/P(t))$, where $P(t)$ are the nominal exchange rates and t is a time label, $t = 1, \dots, N$. Within the rolling window framework we test

	USD/EUR	CAD/EUR	CHF/EUR	GBP/EUR	JPY/EUR
Mean	-4.266e-06	1.715e-06	1.389e-06	-7.607e-06	1.042e-06
Median	0.00000	0.00000	0.00000	0.00000	-1.796e-05
Maximum	0.0181	0.0180	0.0186	0.0174	0.0341
Minimum	-0.0228	-0.0167	-0.0215	-0.0168	-0.0261
Std deviation	1.3e-03	1.4e-03	0.7e-03	1.1e-03	1.7e-03
Skewness	-0.1521	-0.0486	-0.3913	-0.3651	0.2721
Kurtosis	17.59	13.14	52.95	22.03	27.41
JB normality test statistic	333,840	161,300	3,912,000	568,380	934,710
<i>p</i> -value	< 1.0e-03	< 1.0e-03	< 1.0e-03	< 1.0e-03	< 1.0e-03

Table 1: Summary statistics of logarithmic returns

for both linear and non-linear correlations in a window of a specified length n , that is rolling one point forward eliminating the first observation and including the next one, until the last observation. The data in each window are standardized to have zero mean and unit standard deviation, as follows:

$$x(t) = \frac{R(t) - M_R}{S_R}, \quad (1)$$

where M_R and S_R are the sample mean and standard deviation of the window. The window-test procedure [15, 16] uses correlation and bicorrelation potmanteu tests for the detection of linear and non-linear serial dependencies within a time window. The null hypothesis is that the standardized data $x(t)$ in each window are realizations of a stationary pure white noise process. The alternative hypothesis is that the process has some non-zero correlations $C_{xx}(r) = E[x(t)x(t+r)]$ or bicorrelations $C_{xxx}(r, s) = E[x(t)x(t+r)x(t+s)]$ within $0 < r < s \leq L < n$, where L is the number of lags. The sample correlations and bicorrelations in a rolling window are calculated as follows:

$$C_{xx}(r) = \frac{1}{(n-r)} \sum_{t=1}^{n-r} x(t)x(t+r), \quad (2)$$

$$C_{xxx}(r, s) = \frac{1}{(n-s)} \sum_{t=1}^{n-s} x(t)x(t+r)x(t+s). \quad (3)$$

The test statistics in a rolling window for non-zero correlations H_{xx} and b correlations H_{xxx} , are respectively given by:

$$H_{xx} = \sum_{r=1}^L (n-r) C_{xx}^2(r) \sim \chi^2(L) \quad (4)$$

and

$$H_{xxx} = \sum_{s=2}^L \sum_{r=1}^{s-1} (n-s) C_{xxx}^2(r, s) \sim \chi^2((L-1)L/2). \quad (5)$$

We consider several window lengths $n = 2^k$, $k = 3, \dots, 10$. The number of lags is specified as $L = n^b$, where $0 < b < 0.5$ is a user-specified parameter. In order to account only for the presumably most relevant lags, we chose a fixed value of $L = 2$ (i.e. $r = 1, s = 2$), which, for example, for $n = 2^4$ corresponds to $b = 0.25$, as used in Ref. [10].

We note that in order to identify non-linearities by the H_{xxx} test statistics, data pre-whitening is necessary. This is achieved by filtration of the linear component by an autoregressive AR(p) fit. The p order of the AR(p) model is chosen between 1 and $n - 1$ as the optimizer of Schwarz's Bayesian Criterion [26, 27]. The null hypothesis of linear/non-linear correlation is accepted or rejected in each window at a risk level of 5%.

3.2. Simple correlation-based predictor

Above we identify episodes with significant correlations which indicate transient inefficiency of the market. Emergence of such episodes offers some opportunity for the prediction of the price change. In the following we define a simple predictor of the price change direction in the intervals with significant correlations. Prediction of the price change direction \hat{I} (+1 - increase, -1 - decrease, 0 - no prediction) at the time step $t + 1$ is based on the assumption of consistency between the signs of the neighboring price changes and the sample correlation in a window:

$$\hat{I}(t + 1) = \text{sign}(C_{xx}(1)x(t))\theta(\alpha - p_{xx}), \quad (6)$$

where p_{xx} is a p -value of the H_{xx} statistic, θ is the step function ($\theta = 1$ if $p_{xx} < \alpha$ and 0 otherwise), and α determines the confidence level for no correlation null hypothesis rejection (we used $\alpha = 0.05$). Hence, if the sample correlation in a given time window with significant correlation is positive (negative) the predictor predicts the same (opposite) price change direction

as the last observed change. The predictor’s performance is evaluated by a hit rate, i.e. the rate of consistency between the signs of the actual returns and their predictions. We note that we avoided using a more complicated prediction model that would involve parameters. Instead, we tried to link the presence of significant correlations with the degree of predictability as directly as possible, thus eliminating the influence of the parameters’ inference on the prediction. By using the current simple model the only relevant issue is whether the magnitude of the estimated correlation (positive or negative) significantly differs from zero or not.

4. Results

4.1. Frequency and cluster size distribution of significant H_{xx} and H_{xxx} windows

The arrival of periods with significant correlations and their lengths are stochastic processes and their evaluation can shed some light on the overall efficiency of a particular market as well as the interpretation of the events, such as crises, government regulations, etc., that could have led to the increase or decrease of the efficiency [20, 22, 23]. Lim et al. [22, 23] proposed the percentage of time windows in which the market departs from efficiency as an indicator for assessing the relative efficiency. In Fig. 4 we plot the percentage of the significant (a) H_{xx} and (b) H_{xxx} rolling windows for different currency pairs as a function of the window length n . We can observe that the percentage of the significant H_{xx} and H_{xxx} windows tends to increase with increasing n , with some currency pairs (e.g., CHF/EUR) being more and some less (e.g., USD/EUR) sensitive to the variation of n . Based on the efficiency indicator proposed by Lim et al. [22, 23], the largest (smallest) efficiency is found in USD/EUR (CHF/EUR) pair, featuring up to 13.5 (22.2) % of the significant H_{xxx} windows for the considered values of n .

The epochs of the relative inefficiency represent clusters of instants with significant correlations of various sizes. As pointed out in [10, 24], the dependency structures are short-lived and not persistent enough to be properly exploited for prediction purposes. Therefore, in this perspective it is interesting to look into the size distribution of the clusters with significant correlations. In Fig. 5 cumulative distributions are plotted in a log-log scale for the CAD/EUR pair with different values of the window length n . We can see qualitatively similar behavior in both the significant H_{xx} and H_{xxx} cluster size distributions. Namely, for sufficiently large n , fat tails can be

Significant H_{xx} windows				Significant H_{xxx} windows		
n	\hat{x}_{min}	$\hat{\alpha}$	p -value	\hat{x}_{min}	$\hat{\alpha}$	p -value
64	4	1.77	0.00	19	2.82	0.01
128	13	1.97	0.03	19	2.31	0.08
256	12	1.8	0.12	17	1.95	0.14

Table 2: Power-law fits to distributions of clusters of significant H_{xx} and H_{xxx} windows for the CAD/EUR pair for different window lengths n , using the maximum-likelihood fitting method with goodness-of-fit tests based on the Kolmogorov-Smirnov statistic [28]. The parameters \hat{x}_{min} and $\hat{\alpha}$ correspond to the estimates of the lower-bound of the power-law behavior and the scaling exponent, respectively. Statistically significant values with $p > 0.1$ are in bold.

observed. As evidenced from Table 2, a reasonable fit¹ to the power-law behavior, with the estimated exponent $\hat{\alpha} \approx 1.8$ for H_{xx} and $\hat{\alpha} \approx 1.95$ for H_{xxx} , can be obtained by the maximum likelihood method [28] for $n = 256$. This behavior is in clear contrast with the cluster-size distributions of Gaussian white noise, depicted in Fig. 5 by the square symbols. Similar distributions can also be observed using the remaining currency pairs, as shown in Fig. 6.

4.2. Predictability in significant H_{xx} windows

The epochs with significant correlations can be used to forecast the next price change by using the correlation-based predictor (6). Its performance in relation with the significance of the linear correlations is demonstrated in Fig. 7, for a selected period of CAD/EUR data. The cumulative hit rate displays a sharp increase during the epochs characterized by persistent contiguous clusters with very low p -values of the H_{xx} statistic (e.g., the epoch within the green bars). The hit rate typically deteriorates at the end of these epochs due to the fact that, especially for larger window lengths n , the response to changes in the correlation strength is relatively slow and the low p -values persist for some time even after the transient dependences have disappeared. A quicker response can be achieved using smaller window length n , but then the epochs with transient dependences tend to become

¹Normally p -value indicates that the null hypothesis is unlikely to be correct and therefore low values are considered good. In [28], by contrast, the p -value is used as a measure of the hypothesis that is being verified, and hence high values are good. The power law is a plausible hypothesis for the data if the resulting p -value is greater than 0.1. The number of semi-parametric bootstrap repetitions of the fitting procedure is set to 1000.

shorter, as already evidenced in Fig. 5. On the other hand, the predictions made in the epochs with p -values below the threshold $\alpha = 0.05$ but with higher variability are not as successful (e.g., the epoch within the red bars). The tendency of obtaining a better hit rate as the cluster size increases is demonstrated in Fig. 8(a) for the CAD/EUR pair and different values of the window size n . An excellent hit rate is achieved in relatively large clusters for smaller values, such as $n = 8, 16$. However, as shown in Fig. 5, such clusters are scarce and the overall hit rate is governed by the dominant small clusters that show poorer hit rate. For larger values of n (e.g., $n = 256$), the fat tails tell us that large clusters are more ubiquitous, nevertheless, the hit rate in these large clusters is markedly worse than in the those produced by small n . As a result, looking at the total hit rate, no significant differences are observed between different values of the window size n . In Fig. 8(b), we again show qualitative difference compared to the uncorrelated random time series, which does not feature larger clusters with significant correlations and the hit rate is not affected by the cluster size but randomly fluctuates around the value of 0.5.

5. Conclusions

The objective of this paper was to study the dynamics of the linear and non-linear serial dependencies in financial time series, using a windowed portmanteau test procedure in a rolling window framework. In particular, we focused on detection of episodes of statistically significant two- and three-point correlations in returns of several leading currency exchange rates that could offer some potential for their predictability. Since the correlations, which can be viewed as an indicator of market efficiency, are not static but they evolve in time, we employed a rolling window approach in order to capture their dynamics for different window lengths. By analyzing the distributions of the periods with statistically significant correlations it was found that the percentage of the periods with statistically significant two- and three-point correlations in a window of past n observations varies with n and different currencies show different degrees of the dependences. For example, the CHF/EUR pair shows a relatively high percentage of significant correlations, which increases with the window length, while the USD/EUR pair displays a relatively small percentage with little sensitivity to the window length. We found that for sufficiently large window lengths these distributions of the periods with statistically significant correlations fit well to power-law behavior.

We also measured the predictability itself by a hit rate, i.e. the rate of consistency between the signs of the actual returns and their predictions, obtained from a simple correlation-based predictor. It was found that during these relatively brief periods the returns are predictable to a certain degree and the presence of large contiguous clusters of instants with significant correlations increase potential for predictability.

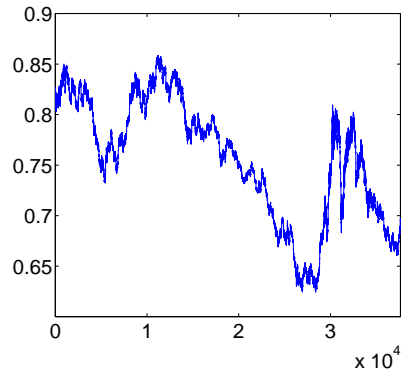
In the present study we did not attempt to use the detected non-linearities in the windows with significant bicomrelations for prediction purposes. Prediction of such non-linear processes is a challenging task, nevertheless, some attempt using a simple model has already been made [29]. In a view of the above, it would be interesting to employ jointly linear and non-linear prediction models during the periods featuring solely linear (non-zero correlation and zero bicomrelation) and non-linear (zero correlation and non-zero bicomrelation) dependencies, respectively.

References

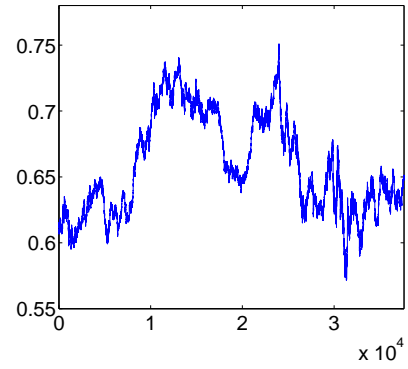
- [1] A. W. Lo, *Econometrica* 59, 1279 (1991)
- [2] C. Eom, S. Choi, G. Oh, W. S. Yung, *Physica A* 387, 4630 (2008)
- [3] C. Eom, G. Oh, W. S. Yung, *Physica A* 387, 5511 (2008)
- [4] H. E. Hurst, *Trans. Am. Soc. Civil Eng.* 116, 770 (1951)
- [5] S. M. Pincus, *Proc. Natl. Acad. Sci. USA* 88, 2297 (1991)
- [6] E. F. Fama, *Journal of Business* 38, 34 (1965)
- [7] C. W. J. Granger, O. Morgenstern, *Kyklos* 16, 1 (1963)
- [8] M. J. Hinich, D. M. Patterson, *Journal of Business and Economic Statistics* 3, 69 (1985)
- [9] K. P. Lim, R. D. Brooks, *Journal of Economic Surveys* 25, 69 (2011)
- [10] A. Serletis, A. G. Malliaris, M. J. Hinich, P. Gogas, *Open Econ Rev*, DOI:10.1007/s11079-010-9194-9
- [11] C. Brooks, M. J. Hinich, *J. Empir Finance* 20, 385 (1999)

- [12] C. A. Bonilla, R. Romero–Meza, M. J. Hinich, *Applied Economics Letters* 13, 195 (2006)
- [13] C. A. Bonilla, R. Romero–Meza, M. J. Hinich, *Applied Economics* 39, 2529 (2007)
- [14] M. J. Hinich, *J Nonparametric Statist* 6, 205 (1996)
- [15] M. J. Hinich, D. M. Patterson, *Mimeo*, University of Texas at Austin (1995)
- [16] M. Hinich, D. M. Patterson, In: M. T. Belongia and J. M. Binner (Eds.), *Money, Measurement and Computation*, (Palgrave Macmillan, London, 2005) 61
- [17] D. O. Cajueiro, B. M. Tabak, *Chaos, Solitons & Fractals* 22, 349 (2004)
- [18] D. O. Cajueiro, B. M. Tabak, *Physica A* 336, 521 (2004)
- [19] J. Alvarez-Ramirez, J. Alvarez, E. Rodriguez, G. Fernandez-Anaya, *Physica A* 387, 6159 (2008)
- [20] Y. Wang, L. Liu, R. Gu, J. Cao, H. Wang, *Physica A* 389, 1635 (2010)
- [21] K. P. Lim, *Physica A* 376, 445 (2007)
- [22] K. P. Lim, R. D. Brooks, J. H. Kim, *International Review of Financial Analysis* 17, 571 (2008)
- [23] K. P. Lim, R. D. Brooks, *Chaos, Solitons and Fractals* 40, 1271 (2009)
- [24] C. Brooks, M. J. Hinich, *Applied Economics Letters* 5, 719 (1998)
- [25] C. M. Jarque, A. K. Bera, *Econ Lett* 6, 255 (1980)
- [26] H. Lutkepohl, *Introduction to Multiple Time Series Analysis*, 2nd edition, (Springer-Verlag, Berlin, 1993)
- [27] A. Neumaier, T. Schneider, *ACM Transactions on Mathematical Software* 27, 58 (2001)
- [28] A. Clauset, C. R. Shalizi, M. E. J. Newman, *SIAM Review* 51, 661 (2009)

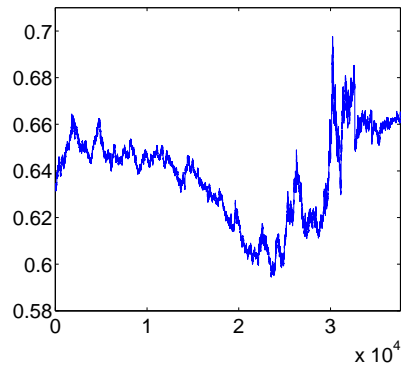
- [29] D. Sornette, V. F. Pisarenko, *Physica D: Nonlinear Phenomena* 237, 429 (2008)



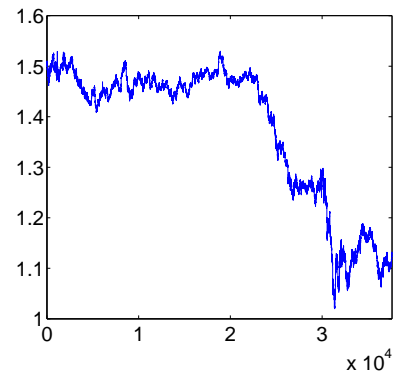
(a) USD/EUR



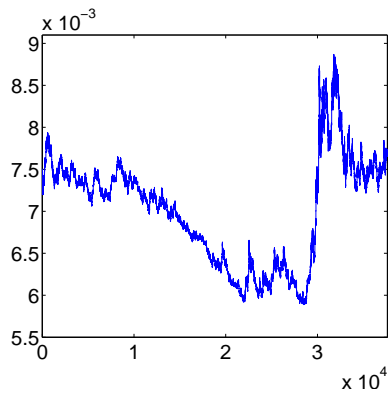
(b) CAD/EUR



(c) CHF/EUR



(d) GBP/EUR



(e) JPY/EUR

Figure 1: Nominal exchange rates.

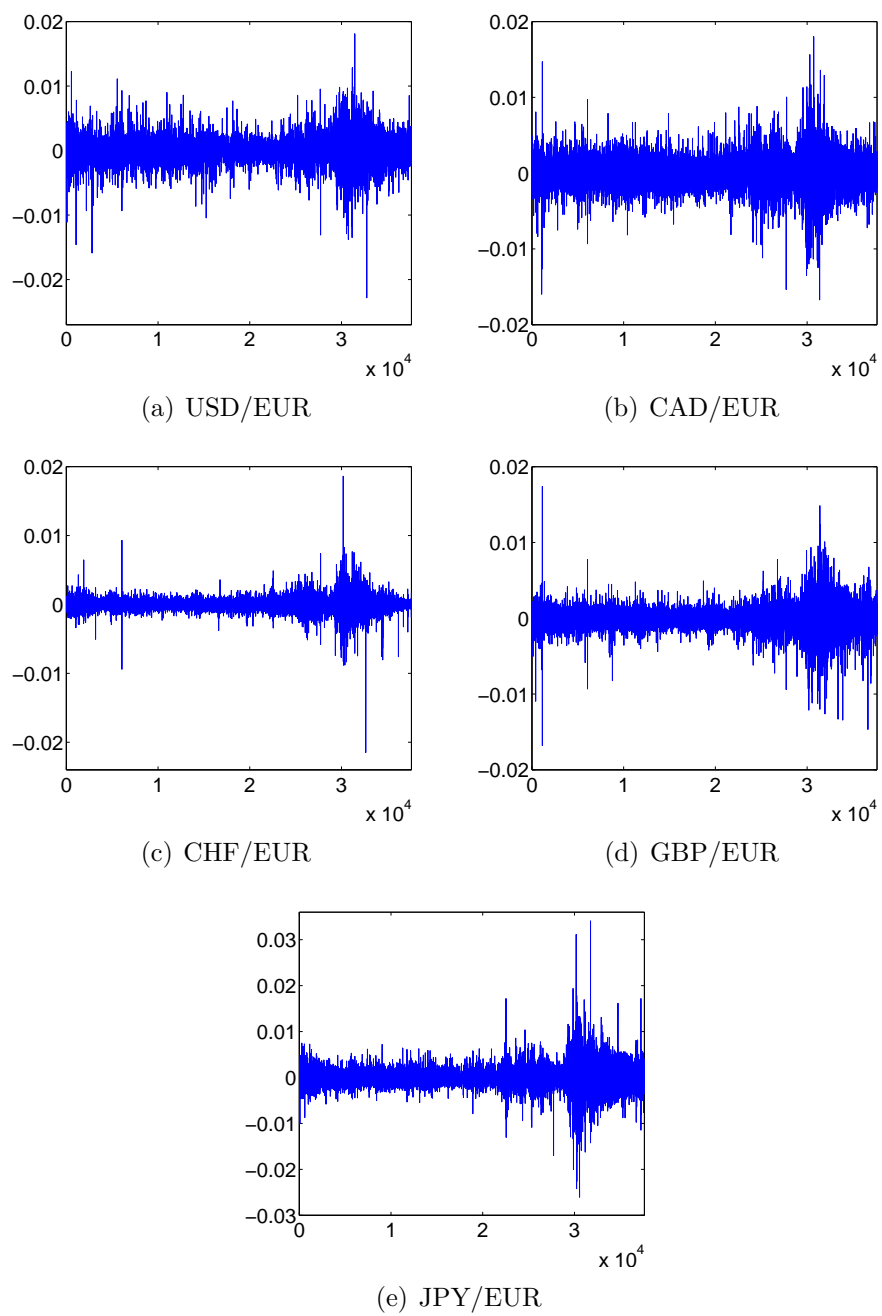


Figure 2: Logarithmic returns.

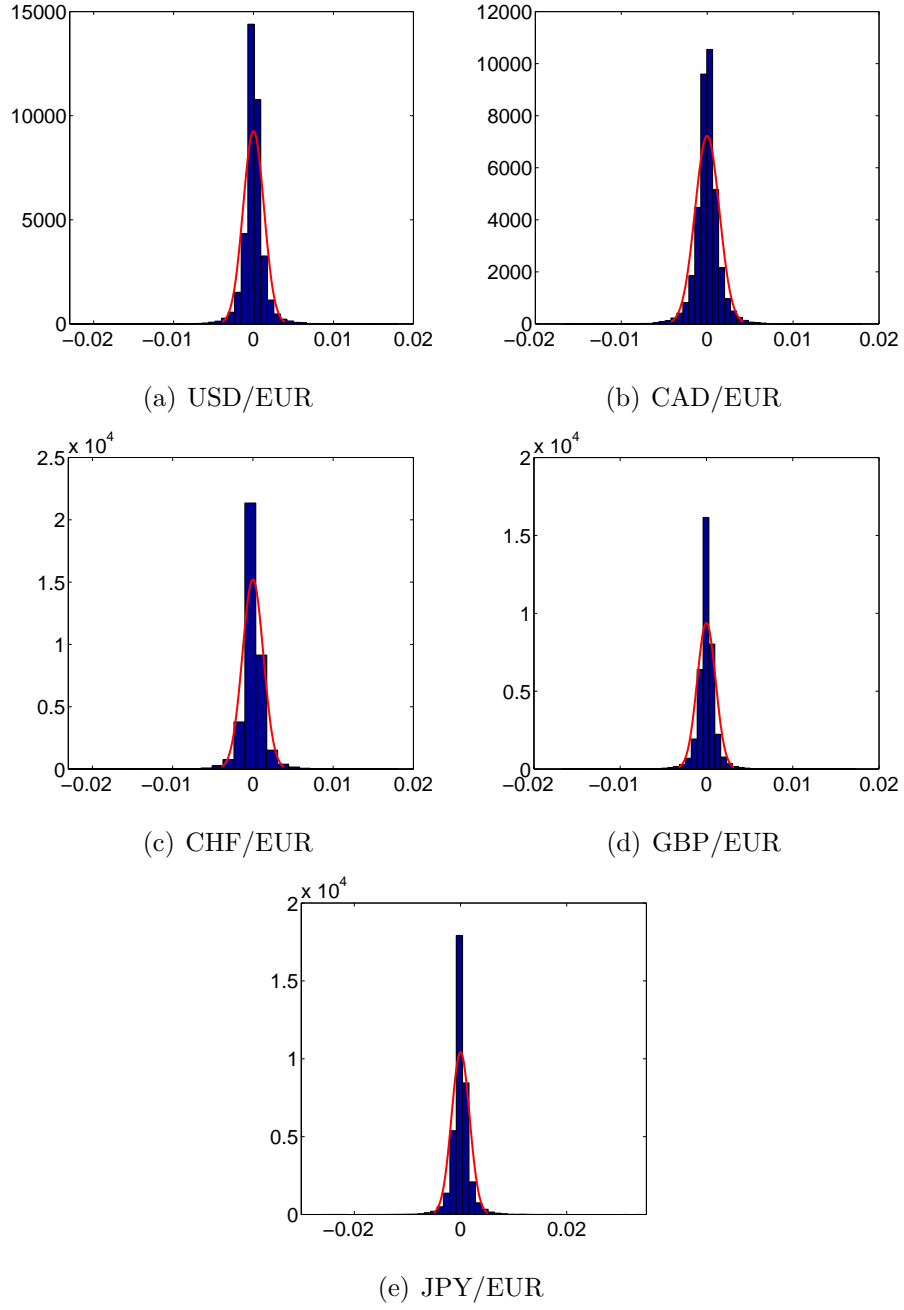


Figure 3: Frequency distributions of the logarithmic returns with the Gaussian curve fit.

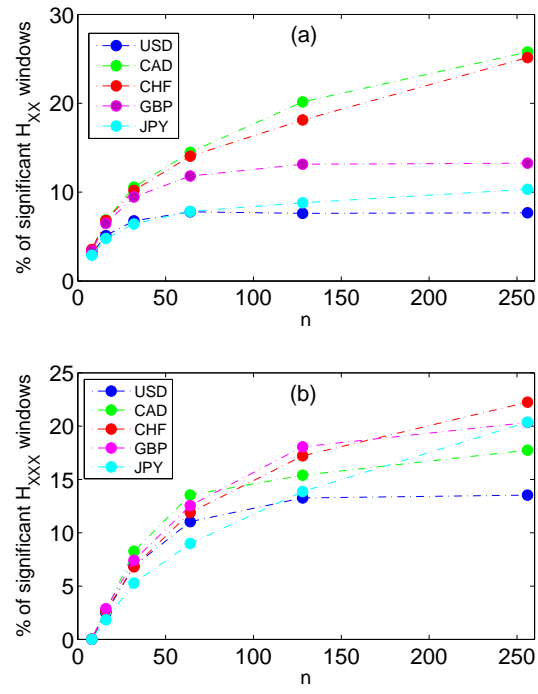


Figure 4: Percentage of the significant (a) H_{xx} and (b) H_{xxx} rolling windows for different currency pairs as a function of the window length n .

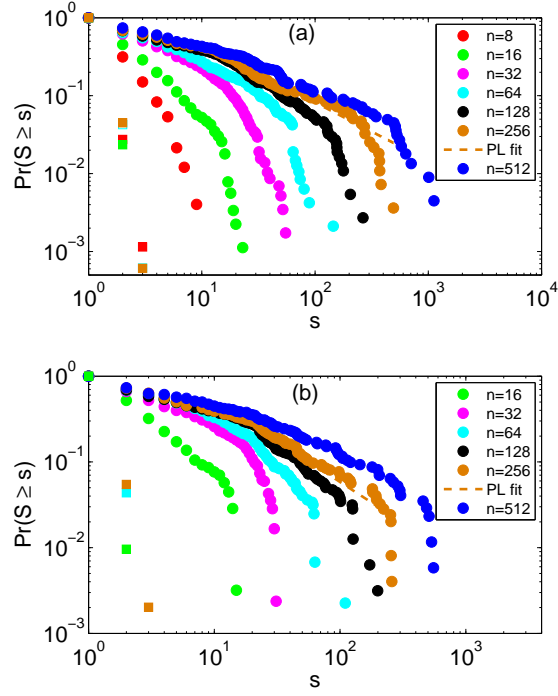


Figure 5: Distributions of clusters of significant (a) H_{xx} and (b) H_{xxx} windows of size s , for the CAD/EUR pair for different values of the window length n . The power-law (PL) fit is done for $n = 256$, using the maximum-likelihood method [28]. The filled squares correspond to a simulated random time series $\sim N(0, 1)$ and different window lengths.

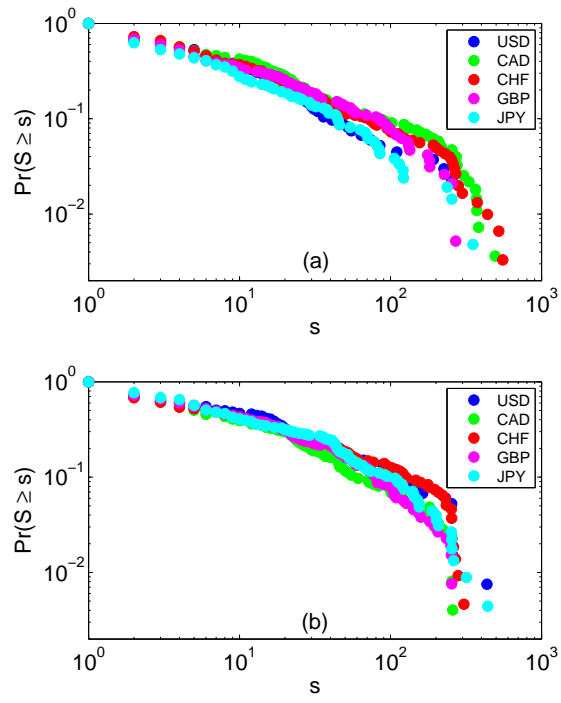


Figure 6: Distributions of significant (a) H_{xx} and (b) H_{xxx} cluster sizes s for different currency pairs for the window length $n = 256$.

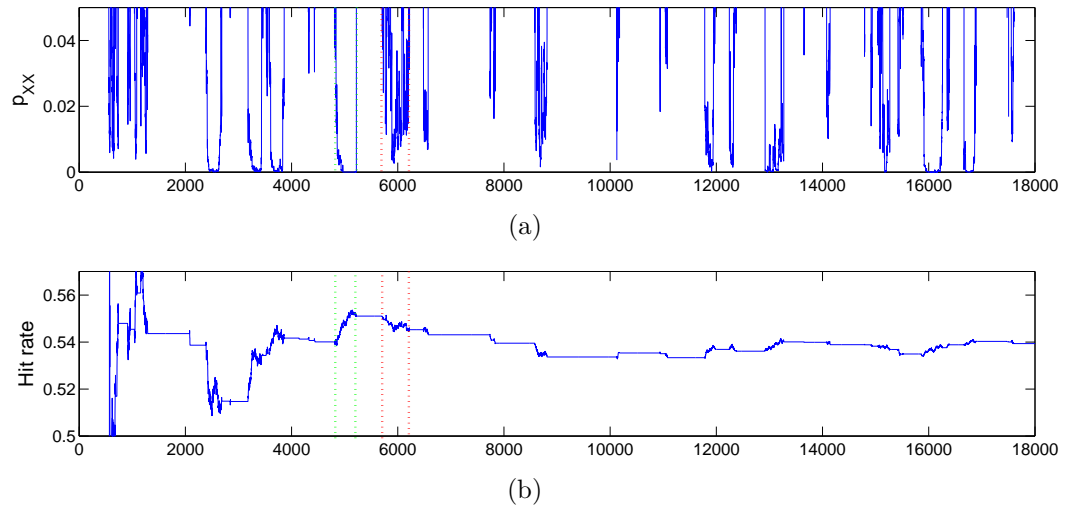


Figure 7: Evolution of (a) the p -value of the H_{xx} statistic (p_{xx}) and (b) the hit rate in the periods of significant correlations. The data correspond to the first 18,000 CAN/EUR return samples and the window size $n = 256$.

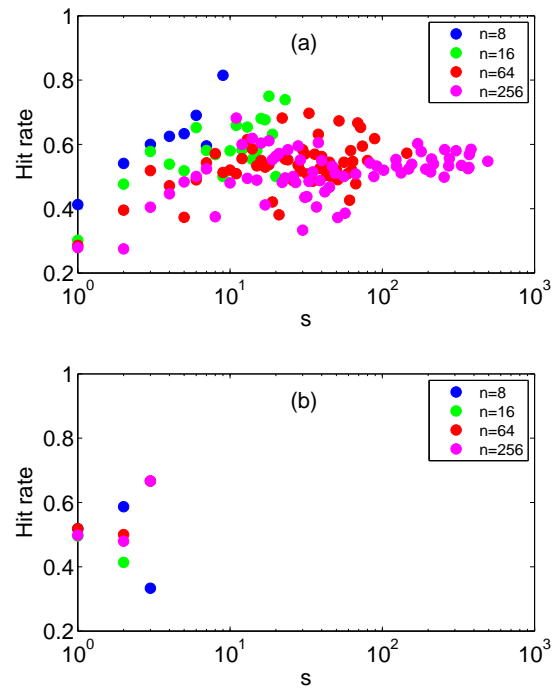


Figure 8: Prediction hit rate as a function of the cluster size s for (a) the CAD/EUR pair and (b) a simulated random time series $\sim N(0, 1)$, using different window lengths n .